

Exhibit I

Self tracking of human motion for virtual reality systems

Dorothy Strickland, Avni Patel, Charles Stovall, Jay Palmer, David McAllister

Department of Electrical and Computer Engineering
North Carolina State University, Raleigh NC 27695

ABSTRACT

Present tracking schemes for virtual reality position sensing have a variety of features that make their use in applications such as large classrooms or remote locations difficult. A more natural tracking method would be a lightweight, low cost, and accurate inertial tracking system. Some commercial inertial systems are discussed. As a low cost alternative, a mouse based head self-tracker has been built at North Carolina State University. Its design and operational ideas are being extended to build a less cumbersome head tracker based on the rotational axes.

1. OVERVIEW OF PRESENT COMMERCIAL TRACKING PRODUCTS

There are three primary philosophies used to track objects in a existing virtual reality (VR) systems¹:

1.1. Free space absolute position sensing

The first involves a device which generates some type of signal and a sensor to receive that signal. Either the signal generator or the sensor is mounted on the person being tracked, and the other device is mounted in the room where motion is to be measured. Often more than one sensor is required for positioning.

This tracking scheme may be implemented with electromagnetic, acoustic, optical, or imaging techniques. While imaging techniques are still in the more experimental stage, the other three systems are available from a variety of vendors. All of these are commonly used with head mounted display (HMD) screens and allow absolute positioning of the user in six full degrees of motion - three spatial and three rotational. A limited variation of this uses a CRT with the same type of sensor mounted on the screen.

Their main common advantage is that the user is allowed to roam somewhat freely in three dimensional space, much as he would if he were actually moving within the virtual reality. This is not always true if the motion detector is mounted in a limited space, such as the front of the computer screen, but there is the potential of reasonable freedom.

Some of the disadvantages at present involve limits in hardware construction, and some involve a more serious problem of the basic philosophy of the design. One construction limit is the allowable range of movement. Although the user can move freely in six degrees within his tracking range, this range is limited in most cases to several meters. A second disadvantage is the limit of motion as the body twists and turns. Bending can cover the position recorder so that it cannot be located by the sensor. Acoustic and optical trackers are especially susceptible to body interference. A third disadvantage is the latency time to determine the positioning accurately, which may include complex mathematical averaging techniques. Another problem is the ergonomics of each design. All are somewhat invasive to the user and some weigh several pounds.

All the problems up to this point deal with one-user difficulties. A limitation of all these systems is that they are designed to track one user in a limited space. Some have included the capability for a limited number of users, but in all cases this involves some trade off in other areas such as range or motion. In addition, all the electromagnetic systems require an extended field outside the tracking range to avoid interference.

1.2. Mounted display screens

A second main technique for displaying the virtual world to the user is to build a floor or ceiling mounted set of screens which can be held up to the eyes by the user, rather than worn as a head-mount. The best known such device, called a boom (binocular omni-oriented monitor) by Fake Labs, contains two CRT's that are viewed through two eye holes. The CRT's are suspended from a mechanical arm that measures the box's position and orientation in space and counterbalances its mass. The user grabs the box by attached handles and moves it around to simulate moving in the scene.

This is much kinder to the neck and back, but more limited in motion range. While the head-mounted systems allow six full degrees of motion, the floor mounted screens can only rotate around the angles of freedom built into the joints of the stand. This does provide a nice range of rotation, but limited spatial movement since the stand is fixed and the user can only move as far as the extensions stretch and only in the angles for which they are designed.

Another advantage is the ease of sensing eye position. Fiber optic feedback from the joint rotation sensors, which are usually encoders, give an accurate positioning for the screens, and thus the eyes. This cuts the two main lag times experienced in the head-mounted displays which result from processing and error correction. A third advantage is that the resolution which can be used with the floor mounted screens is usually much better than the screens that are used in head mounted displays. Most HMDs use LCD's while floor mounted display screens use high resolution CRT's. This is a hardware limitation that is likely to change in the near future.

1.3. Exoskeleton system

The third design, which is much less common for eye tracking, is to use an exoskeleton system for positioning. This concept is a hold over from robotics. Body position is measured from joint angles as the body parts bend. The primary use is in hand positioning.

Most products to date are heavy, expensive, clumsy, and unreliable. The computational problem for compensation for the mass of a control arm can be substantial. Improvements of this idea using fiber optics for joint measurement is going on in several places, including North Carolina State. This method, despite its obvious present problems, has the advantages of not limiting user motion, as in the second method described above, and not confining the user to a specific space. Both of these trade off are important in the self-tracking system described in this paper.

2. SELF TRACKING

2.1. Commercial inertial systems

Inertial tracking systems have been around for some time in military guidance systems. In commercial inertial tracking systems, gyroscopes are combined with accelerometers to provide accurate relative positioning. The moving object is essentially tracking itself locally, relative to its internal directions. Starting with the beginning absolute position in space, the angular motion around three axis is recorded with gyroscopes. Linear motion in the three directions of free space is recorded with accelerometers. By combining these, the full 6 degrees of motion relative to the starting position can be calculated continuously. Since the measurement is relative, rather than absolute, drift may cause errors in the readings which would require an absolute position for re-calibration. Systems, such as the satellite Global Positioning Systems (GPS) are sometimes used for this periodic repositioning.

Such a system would seem ideal for virtual reality, but up to now has not been practical for a variety of reasons. Commercial inertial tracking products have been too heavy and expensive for VR applications. Costs can be hundreds of thousands of dollars. Such systems are physically large, heavy, and require large amounts of power for operation. Solid state design has decreased both the cost and weight of some angular rate sensors (gyro functions) and accelerometers. Systron Donner, a Concord, CA based company has developed a six degrees of freedom inertial

sensor cluster called the MotionPak using its angular rate sensors discussed later in the paper with three servo accelerometers to provide an inertial tracker suitable for human motion. It weighs up to 700 grams, but the components can be broken apart to redistribute the weight. It also has a hefty power requirement of 7 watts, versus 0.8 watts for one of its angular rate sensors. At \$10,000, it is one of the lower cost inertial trackers, and since its price is based on military specifications, it could probably be designed and marketed for considerably less for commercial use.

Accuracy of inertial trackers is still a matter of debate, with often conflicting claims. Even so, these inertial systems are seeing some use in programs where a large range of unrestricted navigation is required, such as in military vehicles, airplanes, and car monitoring. Drift and inaccuracies are updated with map overlays of known points or with a version of the GPS system, although civilian GPS users are assured of only 100 meters navigational accuracy 95% of the time and signals can be blocked by obstacles such as buildings. This would make GPS impractical for re-positioning for human VR motion.

2.2. Advantages of a head motion self tracker

In our research, we are attempting to develop a relative positioning scheme similar to the inertial systems. In our initial design, we have not attempted to provide all six degrees of freedom, but the more limited range of head motion. This will involve the three angular rate changes plus limited linear spatial movement of the eyes around the axis of the head. The tracker is designed to be worn with head-mounted displays.

The trade-offs of such a tracker over the present commercial ones are:

1. Multiple users would be able to operate in the same space without interference. With such a system, you could assign head mounted displays to a class room of individuals and each could exist in their individual virtual realities.
2. Custom rooms would not be necessary. The tracker would be tied to the display helmet, allowing use anywhere there is a connection to a computer and headset for display. This opens the possibility for portable systems, which could be mounted on the user like a backpack.
3. Motion is limited to sitting relatively still with just head action. It would be possible to extend this idea to upper body (or for that matter even full body) motion, but the design complexity would increase greatly. There is some question as to the practicality of such full six degrees of motion in unlimited free space. The physical dangers of moving in a imaginary space while actually walking in a different space would seem to provide insurance liability questions.
4. Drift would eventually create divergence with the actual positioning if relative positioning were used as in true inertial systems. The only practical method of avoiding that is to have some absolute re-calibration mechanism in the room, such as one of the free space absolute methods described above. Such a scheme would remove some of the stated advantages of an independent self tracker by tying the user to a preset room. Re-calibration is not sufficient enough for many VR applications, such as augmented reality, which require that actions be based on absolute position at every point. There is also the question of whether quick, periodic shifts in the virtual image to "correct" it would be more disconcerting to the user than a slow continual misalignment to which the user can adapt.

3. TWO DEGREE OF MOVEMENT HEAD TRACKER DESIGN

Using navigational terminology, the three axes of head motion are defined as: roll, tilting the head to the left or right; pitch, moving the head up or down; and yaw, rotating the head to the left or right. Our tracker is concerned only with roll and pitch, although a method for extending the scheme to include yaw is described. In addition to tracking multi-users and being completely self contained on each user, it was decided the tracker should be light, portable, and low cost. Because there were limited equipment funds, all components needed to be common devices

available in our labs. Although accuracy and latency are major concerns for tracking systems in virtual reality, we did not start with any pre-set conditions; however, these aspects are discussed in the results. This project concentrated more on proving the concept of the tracking idea than in producing a marketable product.

3.1. Initial designs

The initial designs used a variation of commercial tracking systems, with the transmitter and receiver both mounted on the user. One alternative considered was the use of ultrasonic or electromagnetic transmitters mounted on the helmet with the receivers placed on a back plate (Figure 1) or as a sensing coil worn around the neck. This method, however, proposed some of the same kind of problems as the present systems, such as interference signals from nearby users with the electromagnetic tracker. Another design included the mounting of a light weight camera on the helmet and having LED's placed on the back plate. The camera would take pictures of the back plate each time the head moved, and the position of the LED's on the back plate would be different for each picture. Again, this is just a modification of a working optical tracker.




Figure 1. Tracker with back plate.

A variation, which was easier to construct because it did not require the expense of a camera, was to use a mechanical connection between the head and the back plate. A probe was attached to the back of the user's helmet. The probe would possess a ball that moved around on a stationary back plate as the user changed head positions. As a result, each time the head moved, the ball would point to a different position on the back plate. The back plate would consist of touch or pressure pads which would sense the ball and send a signal that would help calculate the angle of the head movement.

By further analyzing the probe design, it was realized that there is a present system that demonstrates the same concept: a computer mouse. If the mouse were placed on the top of the helmet and adjusted so that it could constantly keep contact with the helmet, it would be possible to calculate the head movement, creating a cheap, readily available head tracker.

3.2. Hardware design and construction

Various considerations were what type of mouse to use, the mechanics of the support for the mouse, and helmet size and shape. Each aspect ultimately needed to work together so that the mouse would provide accurate readings with an acceptable resolution.

3.2.1 Choosing the mouse

When considering the type of mouse, the main problem was that of keeping the mouse in contact with a helmet at all times. A regular tabletop mouse was difficult to position around the head curves. As a result, a PC Stylus mouse which is shaped much like a pen was chosen (Figure 2). This stylus shape allowed for easy connection to a mouse support frame.




Figure 2. Stylus and support structure

3.2.2. Designing the support structure

The support structure for the mouse needed to hold the mouse constant for left and right movement, rotate slightly for up and down movement, and be adjustable for various user sizes. Other considerations taken into account were the comfort that it provided the user as well as the size. An aluminum backpack frame that is manufactured by CampTrails was altered for the prototype design. The backpack is lightweight, portable, and allows the user relative freedom of movement.

The actual structure that held the mouse was attached to the backpack. This frame, likewise, was constructed of lightweight aluminum. This structure possesses two rotary joints which allow for the adjusting to different size users (Figure 2). Also, the apparatus moves left to right along the frame of the backpack so that it can be adjusted to the correct calibration point on the helmet. Finally, the mouse was allowed to rotate slightly in the up and down positions so that it would keep constant contact with a helmet at all times.

3.2.3. Designing the helmet

The original plan was to use a basically round, easily available helmet such as a motorcycle, baseball, or football helmet. These types of helmets, however, were analyzed and tested, and it was discovered that the mouse would quickly lose contact with the completely round helmet when the head was tilted down 10.65 degrees. When the head was pitched up, the mouse remained in contact, and when the head was rolled sideways, the mouse lost contact at approximately 40 degrees (Figure 3). Adding a pressure spring to maintain the contact at all times introduced an opposing force which interfered with normal head motion. It was concluded that it was necessary to either re-design the helmet, or construct the mechanical mechanism that supported the mouse so that it possessed constantly moving joints. It was felt re-designing the helmet was the simplest solution.




Figure 3. Illustration of the loss of mouse contact for a round helmet.

For the mouse to remain in contact when the head was pitched down, the helmet must possess a longer and flatter back, similar to a bicycle helmet. Several balloon sizes of 12 inch, 16 inch, and 19 inch were used as reference shapes. They were rotated to various positions with the mouse held stationary. The 19 inch balloon performed best. Six layers of papier mache was used to obtain a rigid structure of the balloon's shape. This structure was placed onto a conventional hard hat that had adjustable neck straps to accommodate different head sizes of various users (Figure 4).

Testing was performed on the prototype as the mouse was held stationary by the backpack rigid frame. The user was successfully able to move his head in all directions without the mouse losing contact. Next, the prototype was tested on several users with different neck sizes and flexibility of rotation. Test subjects ranged in size from a 5 feet 2 inch woman to a six foot man. The mouse remained in contact except in extreme cases, such as looking down more than 75 degrees or tilting left or right to positions where the edge of the helmet made contact with the shoulders. As a result, the helmet design was considered acceptable for this test, and a more rigid prototype was made by covering the original with two fiberglass coatings. In final testing it was observed that the mouse could lose contact if the head motion was too quick, so the final prototype included a rubber mouse pad type coating over the fiberglass.




Figure 4. Front and side final helmet views with mouse

4. CALCULATING EYE POSITION FROM MOUSE MOVEMENT

The output position of the mouse as it rolls along the helmet must be converted to the respective movement of the head. Initially the mouse is placed at a calibration point on the helmet, taken as the zero position. As the head moves, the helmet moves underneath the mouse, changing its output. Since the mouse never leaves the surface of the helmet, its output should in some fashion correspond to the movement of the head in space as well as the rotational angle of the eyes. To determine this relationship, two approaches were tried. The first uses the surface equation of the helmet to determine the head's movement in three dimensions. The second involves using a look-up table of known head positions for various mouse outputs and interpolating for points between those locations.

4.1. Calculating movement from surface shape analysis

Because of the general shape of a 19 inch balloon, the surface was spherical for about half of its length and then became conical towards the rear of the helmet. The fairly primitive construction method of the actual rigid epoxy frame made determining an accurate surface equation from hand measurements difficult. A three dimensional input scanner might have indicated a usable curve, but since one was not readily available, we used the look-up table approach for our tests.

4.2. Calculating movement with a look-up table

The look-up table correlates mouse movement to actual tracked head position and orientation. A specific mouse movement coordinate is obtained in real time as a person moves. If this coordinate matches a value in the table, the tilt and roll eye angles and the x, y, and z spatial offsets for the eye movement are read from the table. If the coordinate point is not in the table, a bilinear interpolation is done using the four nearest table values. This in effect provides an absolute position relative to the mouse's starting calibration point.

To determine the correlation between helmet position and eye position, a measurement grid of one inch squares was placed on the helmet. Using this grid, two independent methods of measuring the look-up table values were performed: video capture and laser pointer.

4.2.1. Picture capture measurement

Using a camera and video digitizer, roll and pitch position for a user at each grid interception was captured. An imaging program was used to print the pictures. For each point on the grid, the user image was captured from the front view to give the roll angle offset and from the side view to give the pitch angle offset. Lines were drawn through the center of rotation and its angle offset from the straight facing calibration point was measured for each grid interception. The size of the pictures and head orientation made it difficult to accurately measure some of the smaller table angles. For that reason, the laser pointer method described below was used to verify the data.

4.2.2. Laser pointer measurement

A pin-point laser was attached on the center of the front-end of the helmet between the eyes and directed at a grid placed on a flat surface. The mouse was positioned at a calibration point on the helmet. The user moved his head until the mouse was at a helmet grid intersection point. The place where the laser was pointing on the flat surface grid was marked to correspond to the helmet grid point. Using geometry and the distance of the user's eyes from the flat grid, the angle of tilt and roll were calculated. There were some problems with this procedure. The lines drawn on the helmet were approximately 0.10 inches thick. This thickness created a point variance of approximately 2 inches for a user 54.4 inches from the measuring plane if different sides of the line were used. Furthermore, it was realized that a slight bending of the aluminum support occurred during certain movement. These problems resulted in a worst case cluster of points for one table measurement where points were up to 1.5 degree offset from the center of the cluster. To account for this, multiple readings were taken for each point to establish the cluster range and the average reading was used to establish each look-up table point.

4.3. Disadvantages of table-look-up method

This table look-up scheme has three inherent errors. The first is the accuracy resulting from the interpolation between points not in the table. A computer program was written to compare the actual movement value with a value derived from linear interpolation for various head angle movements on different slopes of the helmet. The greatest percent difference was found to be only 0.127% for a rotational angle of ten degrees with our helmet shape. Based on these figures, linear interpolation was used and a helmet grid size of 1 square inch was chosen for the table points. A smaller grid would produce better accuracy, but at the cost of larger table size. For our system, the standard mouse interface was used and the look-up table was held in the computer. An extension of the idea might involve a small associative memory look-up table and linear interpolation hardware on the helmet.

Translating the x and y mouse movement into three dimensional space movement is dependent on the individual head arc movement. Although we took measurements with ranges of user size, the data was neither complete nor conclusive enough to give any automatic correlation between such factors as size, sex, and neck length or thickness. NASA studies show some variation in limited studies on joint limits, but again the data is not complete enough for a correlation². It may be that the perception of head position is not critical enough to warrant individual tables. In tests on position of perceived body parts, subjects dissociated actual and perceived position of the head with respect to the body. For example, in one test, subjects with their eyes closed and their head turned for a long time begin to feel that the head was slowly returning to the neutral position. After 10 to 12 minutes of sitting with the head turned, the error in perceived head position reached up to 60 degrees³. We are continuing tests to determine if it is possible or necessary to generate automatic data from some combination of user physical structure.

There exists a dead spot in the axis around the mouse where head motion is not sensed as movement. The calibration point was taken off center on the side of the helmet to make this spot less conspicuous. The solution to

this problem is to add a second mouse on another part of the helmet. This would also pick up the third axis, providing yaw (side-to-side) motion sensing. The two readings taken together would provide full head positioning.

5. TRACKER SOFTWARE

The software part of the tracking system translated the mouse's output into the user's head position using the look-up table. The repeating general algorithm is as follows: 1) get the mouse's movement from the last position; 2) add or subtract this amount to the total distance from the starting point; 3) use this new total to find the nearest grid points in the look-up table; and 4) use these values or interpolate (if necessary) to determine the actual head position.

Communication with the mouse and the mouse driver was done with the use of a special interrupt. The standard Microsoft supplied mouse driver was used. According to the vendor specifications, the distance counters, which can be positive or negative, are only reset after being read. Our testing indicated problems with the Microsoft provided motion function which resulted in the motion counters losing counts at high head acceleration. According to the vendor, this problem is being corrected. The Stylus mouse used has a hardware resolution of 1/400 of an inch. The code to interpret the input and interpolate to the appropriate table values was written in C. The operating system functions were reduced to a minimum and all other applications were suspended while testing. The resulting execution time for one cycle of the tracker was 0.0333 seconds on an Intel 80386 processor running at 16 MHz. This is within published acceptable latencies assuming the computer can display the simulated images for the new tracker location in 0.0666 msec. Although the FAA latency requirements for commercial flight simulators (level C) is 150 msec⁴, studies have shown that user perceived visual quality degrades steadily when delays in time from motion to visual presentation is over 100 msec⁵.

6. SUMMARY OF DESIGN

The system was tested by having a user move his head in a normal manner and observing the tracked position read-out on the computer CRT display. The correctness of a point was verified by stopping motion and measuring the actual parameters versus the displayed ones. For normal head motion speeds, the system was found to be reasonably accurate. The error ranges and problems encountered are described above in relevant design component sections. For a test of feasibility of a self-tracker, it was felt the mouse design was successful. The discomfort of the helmet and mouse support structure led us to take this basic idea of motion displacement from pitch and tilt and extend it to a more natural scheme using sensors that measure axis movement.

7. EXTENSION USING AXIS MOVEMENT SENSORS

Using the idea verified with the mouse that head arc motion follows a repeating algorithm, it should be possible to calculate the user head position with measurements of movement around the three axes. Assuming a consistent, measurable head axis arc, each combination of tilt, pitch, and yaw would produce one distinct head orientation and spatial position for an individual. Both tilt sensors and gyroscopes were investigated as rotational angle sensors.

Modern gyroscope appear to be the more sophisticated of the two. A piezoelectric quartz angular rate sensor produced by Systron Donner of Concord, CA weighs approximately 30 grams and has a resolution of 0.002 degrees/sec for motion in the 100 degree/second range. It is capable of handling movement up to 1000 degrees/sec, an acceptable maximum for most head motion. Typical drift of the sensor after stabilization is reported to be 5 degrees per hour. We tested this sensor, called a GyroChip, by mounting it on a helmet and attaching it to a Intel 80386 machine to give a rotational movement along one axis. Three such chips would provide accurate readings of all three axes. Disadvantages were that it required a low power supply and A/D converter to use standard input into our PC. This increased the weight of the tracker. However, our main deterrent on using it for further tests was its cost, about \$1,000 per GyroChip in quantities of 10-50.

Another rate sensor is produced by Gyration, a Saratoga CA company. The 1.2 oz GyroEngine uses an LED and an optical sensor to provide angular accuracy of 0.1 degree for normal head motion. A representative of the company indicated it could handle head accelerations up to 1000 degrees/second. Published specifications indicate it may have more drift than the GyroChip. It is still rather costly at \$500 per gyro, but only two are necessary as each measures two axes of motion. It also requires a low power supply but the output is digital, eliminating the need for an A/D converter.

Murata Products and Matsushita Inc. produce vibrating beam based gyroscopes that are also small and use low power. These have been used in automotive navigation systems in competition with the GyroEngine.

We are presently developing a fluid sensor based on measuring the capacitive difference generated as fluid moves in a tube. Such a sensor could hopefully be constructed for under \$10. It will, however, only measure tilt and pitch, not yaw.

8. CONCLUSIONS

With present commercial development of inertial trackers, it may not be long before a practical, self-tracking system is available for virtual reality systems. Such systems, however, will still be relatively expensive in the near future. It would be useful to have an inexpensive tracker which would not limit the user to a particular position in a preset room. The mouse based tracker shows the potential for a self-tracker using rotational axis movement to determine orientation and spatial positioning. Extending this idea to modern gyro and sensor technology could provide the user with a truly portable head tracker. This would open up new applications which are not limited to a graphics lab. Presently we are developing this tracking method for use with a virtual reality application to help treat children with physical and mental disorders, such as autism and attention deficit hypersensitivity disorder, in the local school system.

9. ACKNOWLEDGMENTS

The authors wish to thank Norman Badler and Francisco Azuola of Penn State for sharing the NASA joint limb data, and Systron Donner for their loan of the GyroChip for evaluation.

10. REFERENCES

1. L. Harrison and D. McAllister, "Implementation issues in interactive stereo systems," Stereo Computer Graphics and Other True 3D Technologies, edited by D. F. McAllister, pp. 119-151, Princeton University Press, Princeton, New York, 1993.
2. "Joint Limits Data-NASA", NASA Man Systems Integral Manual::1987.
3. V. S. Gurfunkel and Y. S. Levick, "Perceptual and automatic aspects of the postural body scheme," Brain and Space, edited by J. Paillard, Oxford University Press, pp. 147-162, 1991.
4. Airline Simulator Qualification, Advisory Circular 120-40B, Federal Aviation Administration, May 1989.
5. P. Beckett, "Effective cueing during approach and touchdown-comparison with flight," In Piloted Simulation Effectiveness of AGARD Conference Proceedings 513, Brussels, 14-17, pp. 30.1-30.11, Oct. 1991.